

On the Rejectability of the Subjective Expected Utility Theory*

Konrad Grabiszewski[†]

February 2012

Abstract

State space is an element of the Subjective Expected Utility (SEU) theory that is constructed in the agent's mind but is not directly observable. The researcher who verifies whether or not the agent violates the SEU theory could presume a state space but he risks reaching conclusions based on false assumptions. As an alternative approach, I propose SEU-rationalization: From the agent's observable choices, the researcher constructs the state space and belief over that state space such that the agent appears to satisfy the SEU theory. I derive conditions under which it is possible to SEU-rationalize the agent's behavior.

Keywords: subjective expected utility theory; state space; SEU-rationalization

JEL classification: D01; D03; D81.

*I would like to thank Itzhak Gilboa, Alex Horenstein, and Erkut Ozbay for their valuable comments. Support from the Asociación Mexicana de Cultura A.C. is gratefully acknowledged.

[†]Address: Department of Business Administration, ITAM, Av. Camino a Sta. Teresa No. 930, Col. Magdalena Contreras, C.P. 10700 Mexico D.F., Mexico, konrad.grabiszewski@itam.mx, <http://allman.rhon.itam.mx/~konrad.grabiszewski/>

1 Introduction

An American firm plans an expansion to the European market. Among other things, the managers must decide where to locate their European headquarters. How can they solve this problem? If the firm chooses to rely on the Subjective Expected Utility (SEU) theory developed by Leonard J. Savage (Savage (1972)), then the first step is preparation of the list of all feasible options. Such a list constitutes the **set of alternatives** and is denoted by $\mathcal{F}$. Let f stand for France, g for Germany, h for Hungary, and so on. The second step then involves constructing **state space**, as denoted by Ω . Each state, ω , is then, in the words of Savage, “a description of the world, leaving no relevant aspect unprescribed.” In the case of the location problem, there could be several political, social, or economic uncertainties whose realizations are not known before the firm makes a decision, though they affect the firm’s business objectives. For instance, France can change its 35-hour work week to a 30-hour work week, or Germany can decide to withdraw from the euro zone, or Hungary can experience a crisis as severe as the ongoing Greek debt crisis. Each given state represents a realization of all uncertainties that the firm deems relevant. For example, if state ω_1 is true, then France does not change its labor law, Germany remains in the euro zone, and Hungary suffers a huge crisis. It is assumed that states are internally consistent (i.e., no ω represents a scenario in which France decreases and, simultaneously, increases the maximum of weekly working hours) and only one state is true, though the firm is uncertain which state is/will be true. The next element of the SEU theory is the **set of consequences**, X . The consequences, also called outcomes, are what the firm ultimately cares about. They can be expressed as monetary gains, but this is not necessary. For example, the firm can consider a three-element set of consequences, $X = \{x_1 = \text{“very good business situation,” } x_2 = \text{“acceptable business situation,” } x_3 = \text{“bad business situation”}\}$. For a given state, each alternative yields a specific consequence. For instance, if ω_1 is true, then the outcomes of choosing f , g , and h are x_1 , x_1 , and x_3 , respectively. Using the state space and set of consequences, the firm represents alternatives as functions, also known as acts,

$$f : \Omega \rightarrow X. \tag{1}$$

As such, the set of alternatives is called the **set of acts**. Using the state space, the set of consequences, and the set of acts, the firm is able to construct a utility function, u , defined on X , $u : X \rightarrow \mathbb{R}$, and a subjective probability measure, λ , defined on Ω such that each act is evaluated by its subjective expected utility,

$$\int_{\Omega} u(f(\omega)) d\lambda. \tag{2}$$

The firm then picks a location that yields the highest expected utility.

This summarizes how SEU can be used to solve decision problems. In this paper, however, I do not analyze the normative, but rather the descriptive, aspects of the SEU theory. That is, I focus on the problem of the **researcher** (hereafter, referred to as “he”) who tries to determine whether or not the behavior of the **agent** (hereafter, referred to as “she”) violates the SEU theory.

From the researcher’s perspective, the key problem with testing the SEU theory is that he can tell *what* choices the agent made, but in general he does not know *how* she made these choices. In the case of the American firm, the researcher does not participate in the business meetings and does not monitor the firm’s analysis of alternatives. Hence, the researcher does not know what uncertainties the managers consider to be relevant and, in consequence, what state space the firm may be considering. It is the non-observability of state space that is a focus of this paper. In order to test the SEU theory, the researcher, who has only a limited understanding of how the agent analyzes her decision problem, could make assumptions about the agent’s analysis—for example, he could assume a specific state space. However, in a such case the researcher is bound to make incorrect conclusions since his assumptions could be false.

In this paper, I propose an alternative approach in which I consider the most prudent researcher, who does not presume any state space. This researcher relies on data consisting of observable choices and uses it to construct a state space and belief over that state space such that the agent appears to satisfy the SEU theory. I call this procedure **SEU-rationalization** of the agent’s behavior. As part of this, I assume that the researcher (a) knows what set of alternatives the agent faces; (b) observes the agent’s **preference relation**, $\succsim$, defined over the set of alternatives; (c) knows what set of consequences the agent considers; and (d) knows how the agent perceives the alternatives. Since there is no state space, the alternatives are translated into subsets of the set of consequences, X , instead of acts. That is, an alternative, f , is identified with a subset of X , X_f , which is interpreted as the collection of all possible consequences that the alternative can yield.

The main question of this paper is:

Under what conditions is it possible to SEU-rationalize the agent’s behavior?

I use the following two-step analysis to obtain the answer.

Step 1. In Section 3, Lemma 3.1 develops conditions characterizing the set of alternatives, the agent’s preference relation, and the set of consequences that imply the existence of a utility function, u , defined over the set of consequences, X , as well as, for each alternative, f , a probability measure, λ_f , defined over X such that the support of λ_f is X_f and f is evaluated by its expected utility, $V(f) = \int_{X_f} u(x)d\lambda_f$.

Step 2. In Section 4, Proposition 4.1 constructs a state space, Ω , and a probability measure, λ , defined over that state space such that the alternatives become acts that are evaluated by their expected utility, $V(f) = \int_{\Omega} u(f(\omega))d\mu$. The construction of Ω and λ relies on X_f s and λ_f s from Lemma 3.1.

This paper is not the first to recognize that the researcher does not directly observe states and that this lack of observability make it difficult to verify whether or not the agent violates a given theory. For instance, Gilboa (2009) notes that “when we consider a real-life choice, different people may think of different eventualities, different possible outcomes, and plan for different time horizons. Actual choices may not satisfy P2 [i.e., an axiom of Savage] when embedded in a given model, and yet they may satisfy it when embedded in a larger model.” Similar observations are made in Epstein (2010), Karni (2008), and Machina (2003). This paper

adds a method of SEU-rationalization to the existing work. An example of applying this method is presented in Section 2, which includes two SEU-rationalizations of the Ellsberg three-color urn experiment—hereafter called the Ellsberg experiment (Ellsberg (1961)). One of these SEU-rationalizations implies that the Ellsberg experiment does not constitute a valid empirical argument against the SEU theory.

This paper is constructed in the following way. Section 2 contains two SEU-rationalizations of the Ellsberg experiment. Sections 3 and 4 develop a method that allows to conduct an SEU-rationalization. Section 5 discusses the place of this paper in the literature, and Section 6 contains final remarks. All proofs are relegated to the Appendix.

2 SEU-Rationalization of the Ellsberg experiment

In this section, I present a very simple version of the Ellsberg experiment and show that the experiment’s result can be SEU-rationalized when the agent’s behavior is allowed to be irrational.

In the urn, there are three balls. One of them, and only one, is red. Each of the remaining two can be either green or blue. However, the number of green balls in the urn is unknown. One ball will be drawn from the urn. The agent has two decisions to make. The first choice is between a bet on a red ball, f_r , and a bet on a green ball, f_g . The second choice is between a bet on a red or blue ball, f_{rb} , and a bet on a green or blue ball, f_{gb} . Without loss of generality, I assume that betting successfully yields a utility of 1 while betting unsuccessfully generates a utility of 0. The following preference relation is of interest:

$$f_r \succ f_g \tag{3}$$

$$f_{gb} \succ f_{rb}$$

The set of consequences, X , consists of two elements, S (uccess) and F (ailure). How can the agent construct a state space to analyze this problem? Consider $\hat{\Omega} = \{R, G, B\}$, where the letter indicates which ball will be drawn. This can be called a posterior-view state space, since the states represent what will be true after the drawing occurs. Given $\hat{\Omega}$, it is possible to represent alternatives as the Savage acts—that is, the functions from $\hat{\Omega}$ to X . For instance, $f_r(R) = S$ and $f_r(B) = f_r(G) = F$. Let $\hat{\lambda}$ be a probability measure on $\hat{\Omega}$. The preference relation depicted in (3) is inconsistent with the SEU theory: $f_r \succ f_g$ implies that $\hat{\lambda}(\{R\}) > \hat{\lambda}(\{G\})$ while $f_{gb} \succ f_{rb}$ implies that $\hat{\lambda}(\{G, B\}) > \hat{\lambda}(\{R, B\})$, which indicates that $\hat{\lambda}(\{G\}) > \hat{\lambda}(\{R\})$. Another obvious choice of state space is $\tilde{\Omega} = \{0, 1, 2\}$, where each state stands for the number of green balls in the urn. This can be called a prior-view state space, since states represent what is true before the draw. In this context, bets are maps from $\tilde{\Omega}$ to the set of (objective) probability measures on X , $P(X)$, which implies the use of Anscombe and Aumann (1963). It is straightforward to show that the Ellsberg experiment also violates the Anscombe-Aumann theory.

Savage's model with state space $\hat{\Omega}$ does not explain (3), but allowing for another state space, it might be possible that there is a Savage model that generates (3). Table 1 describes such an SEU model. (Section 4 provides the details of how Table 1 was constructed.)

state	f_r	f_g	f_{gb}	f_{rb}	probability	state	f_r	f_g	f_{gb}	f_{rb}	probability
ω_1	S	S	S	S	120/1296	ω_9	F	S	S	S	60/1296
ω_2	S	S	S	F	120/1296	ω_{10}	F	S	S	F	60/1296
ω_3	S	S	F	S	24/1296	ω_{11}	F	S	F	S	12/1296
ω_4	S	S	F	F	24/1296	ω_{12}	F	S	F	F	12/1296
ω_5	S	F	S	S	240/1296	ω_{13}	F	F	S	S	120/1296
ω_6	S	F	S	F	240/1296	ω_{14}	F	F	S	F	120/1296
ω_7	S	F	F	S	48/1296	ω_{15}	F	F	F	S	24/1296
ω_8	S	F	F	F	48/1296	ω_{16}	F	F	F	F	24/1296

Table 1: SEU-rationalization of the Ellsberg experiment.

In Table 1, there are sixteen rows, with each denoting a state. Columns f_r , f_g , f_{gb} , and f_{rb} show the realizations of bets at a given state. For example, at ω_1 all bets yield S . The probability column depicts the probability that a given state will occur. If $V(f)$ denotes the expected utility of act f , then it is easy to verify that the SEU model in Table 1 captures the choices in (3):

$$V(f_r) = \frac{120 + 120 + 24 + 24 + 240 + 240 + 48 + 48}{1296} = \frac{864}{1296}$$

$$V(f_g) = \frac{120 + 120 + 24 + 24 + 60 + 60 + 12 + 12}{1296} = \frac{432}{1296}$$

$$V(f_{gb}) = \frac{120 + 120 + 240 + 240 + 60 + 60 + 120 + 120}{1296} = \frac{1080}{1296}$$

$$V(f_{rb}) = \frac{120 + 24 + 240 + 48 + 60 + 12 + 120 + 24}{1296} = \frac{648}{1296}$$

The main difference between state space as given in Table 1 and the previously mentioned $\hat{\Omega}$ is that the former allows for irrationality. For instance, at ω_5 only act f_g does not yield success—that is, the drawn ball has two colors, red and blue. Other states are even more baffling—for example, at ω_8 the drawn ball is

red and, at the same time, not red. The SEU-rationalization in Table 1 relies on the objectively irrational probability measure. The notion of objective and subjective rationality was established by Gilboa et al. (2010). The agent's belief is objectively rational if she is able to convince others that her belief makes sense. If others cannot convince the agent that her belief is wrong, then such a belief is subjectively rational. In the Ellsberg experiment, it is objectively rational that, for example, if a red ball is drawn then the probability of the same ball being blue (or, green) or not being red is 0. Thus, ω_5 and ω_8 should be assigned a probability of 0. In fact, the only states that are objectively rational are ω_7 (i.e., red ball is drawn), ω_{10} (i.e., green ball is drawn), and ω_{13} (i.e., blue ball is drawn). These are the states R , G , and B , respectively, from the state space $\hat{\Omega}$.

In Table 1, the agent behaves as if she is assigning a probability of $\frac{1068}{1296}$ to the set of states that are objectively irrational. Measuring the agent's irrationality by the weight she assigns to such states, it is fair to say that the SEU-rationalization in Table 1 is not satisfactory, since it relies on large irrationality. Therefore, it is obvious to ask: What is the smallest irrationality that would still allow for the SEU-rationalization of the Ellsberg experiment? The answer is: As small as desired. Let ε be a total mass assigned to objectively irrational states. Fix ε as smaller than 0.3 and consider the SEU model given in Table 2.

state	f_r	f_g	f_{gb}	f_{rb}	probability	state	f_r	f_g	f_{gb}	f_{rb}	probability
ω_1	S	S	S	S	0	ω_9	F	S	S	S	0
ω_2	S	S	S	F	0	ω_{10}	F	S	S	F	$\frac{1}{10} - \frac{1}{3}\varepsilon$
ω_3	S	S	F	S	0	ω_{11}	F	S	F	S	0
ω_4	S	S	F	F	0	ω_{12}	F	S	F	F	0
ω_5	S	F	S	S	0	ω_{13}	F	F	S	S	$\frac{8}{10} - \frac{1}{3}\varepsilon$
ω_6	S	F	S	F	ε	ω_{14}	F	F	S	F	0
ω_7	S	F	F	S	$\frac{1}{10} - \frac{1}{3}\varepsilon$	ω_{15}	F	F	F	S	0
ω_8	S	F	F	F	0	ω_{16}	F	F	F	F	0

Table 2: ε -small irrationality and SEU-rationalization of the Ellsberg experiment.

Next verify that the SEU model depicted in Table 2 captures the choices in (3):

$$V(f_r) = \varepsilon + \frac{1}{10} - \frac{1}{3}\varepsilon = \frac{1}{10} + \frac{2}{3}\varepsilon,$$

$$V(f_g) = \frac{1}{10} - \frac{1}{3}\varepsilon,$$

$$V(f_{gb}) = \varepsilon + \frac{1}{10} - \frac{1}{3}\varepsilon + \frac{8}{10} - \frac{1}{3}\varepsilon = \frac{9}{10} + \frac{1}{3}\varepsilon,$$

$$V(f_{rb}) = \frac{1}{10} - \frac{1}{3}\varepsilon + \frac{8}{10} - \frac{1}{3}\varepsilon = \frac{9}{10} - \frac{2}{3}\varepsilon.$$

Explanations for the behavior of the Ellsberg experiment's subjects can be divided into two camps.

- (a) The non-SEU explanations rely on assumption that the agent is objectively fully rational—that is, her state space is $\{\omega_7, \omega_{10}, \omega_{13}\}$. The non-SEU explanations consists of the ambiguity theories. The ambiguity literature begins with Schmeidler (1989), who models the agent's belief as a non-additive measure. As an alternative, in Gilboa and Schmeidler (1989), the agent is characterized by the set of priors and acts are evaluated by the minimal expected utility they yield. In Ghirardato et al. (2004), which extends the model of Gilboa and Schmeidler (1989) the agent ranks acts according to the weighted sum of their minimal and maximal expected utilities. Another generalization of Gilboa and Schmeidler (1989) is provided by Maccheroni et al. (2006). In their model of variational preferences, the agent takes into account not only the minimal expected utility that an act generates but also an ambiguity index. One special case of the model of variational preferences is the multiplier preferences model introduced by Hansen and Sargent (2001) and axiomatized by Strzalecki (2010). Yet another non-SEU approach involves modeling the agent as if having second-order belief. In Segal (1987), the agent's decision-making process consists of two stages in which the agent constructs a belief over the state space, as well as the set of probabilities over the state space. Finally, Siniscalchi (2009) offers a model of vector expected utility in which acts are the sum of subjective expected utility and adjusting function. (See Eichberger and Kelsey (2009), Ghirardato (2010), and Siniscalchi (2008) for a review of the ambiguity literature. The most complete and up-to-date bibliography is Wakker (2011).)
- (b) The SEU-based explanations enlarge the state space, $\{\omega_7, \omega_{10}, \omega_{13}\}$, and require the agent to assign non-zero weight to that enlargement. Gilboa and Schmeidler (1994) and Mukerji (1997) belong to this group as well as the SEU-rationalizations in Tables 1 and 2 do. Due to the large weight it assigns to irrational states, I omit the SEU-rationalization in Table 1 from further discussion. In Table 2 the total mass assigned to irrational states is ε . In fact, there is only one irrational state, ω_6 , that has non-zero weight. The existence of such irrationality can be due to the agent not being accustomed to decision problems expressed in terms of the urn and balls. As Gilboa (2009) writes, "David Schmeidler often says, 'Real life is not about balls and urns.' Indeed, important decisions involve war and peace,

recessions and booms, diseases and cures.” When faced with the problem described in an unfamiliar fashion, the agent might misunderstand the problem or miscalculate the probability, and make an ε -small mistake. Alternatively, in the laboratory environment, the ε -small irrationality might be due to the agent mistrusting the experimenter—that is, the agent believes that she is playing a game with the experimenter in which the experimenter has superior knowledge about the urn or is trying to manipulate her. This line of thought is not new and has been pursued by, among others, Kadane (1992) who shows that the two-urn Ellsberg experiment fits within the SEU framework as long as the agent assigns arbitrarily small probability to the experimenter being malevolent. (See also Brewer (1963), Frisch and Baron (1988), Hey et al. (2010), Kühberger and Perner (2003), Morris (1997), Mukerji and Shin (2002), Pulford (2009), and Trautmann et al. (2008).)

From an empirical perspective—that is, from a perspective that bases conclusions only on observable data—the fact that only ε -small irrationality is enough to SEU-rationalize the Ellsberg experiment implies that the experiment does not reject the SEU theory. This claim is based on two arguments. First, suppose that the only available observations are those in (3). Since the state space is unobservable, there is no empirical reason to assume (or not to assume) that the state space is $\{\omega_7, \omega_{10}, \omega_{13}\}$ (i.e., non-SEU explanations) or the enlarged $\{\omega_7, \omega_{10}, \omega_{13}\}$ (i.e., SEU explanations). In this case, the data reject neither of these explanations. Second, suppose that the researcher does have more information about the agent’s choices and is able to estimate the weight that the agent’s belief assigns to the enlargement of $\{\omega_7, \omega_{10}, \omega_{13}\}$. In the case of the SEU-rationalization given in Table 2, that enlargement consists of state ω_6 . The only difference between the non-SEU explanations and this given in Table 2 is the weight assigned to state ω_6 : 0 and ε , respectively. However, statistically speaking, 0 and ε are the same—that is, it is not possible to reject the hypothesis that the agent assigns weight ε to ω_6 in favor of the hypothesis that the weight is 0 (or, vice versa). In situations like this, Gilboa and Samuelson (2011) indicate that “[p]eople typically bring subjective criteria to bear in making this choice [i.e., choosing among theories], tending to select theories that seem a priori reasonable, intuitive, simple, elegant, familiar, or that satisfy a variety of other considerations.” In the case of the Ellsberg experiment, it is left for the reader’s subjective criteria to decide which explanation seems more appropriate. One such a non-empirical criterion that can be used to bolster the position that the Ellsberg experiment is a valid argument against the SEU theory is the requirement that the agent be fully rational.

3 Step 1: Uncertainty without state space

3.1 Terminology and notation

There is an agent and there is a researcher who analyzes her. The agent faces the set of alternatives denoted by $\mathcal{F}$, and I assume that the researcher knows that the set of alternatives is $\mathcal{F}$ and not some other set. (Note that here and later, “to know” has an informal meaning and is not formally defined as it is in the epistemic literature.) The agent is characterized by a preference relation, $\succsim$, defined over $\mathcal{F}$, and I assume that, relying on the agent’s observable choices, the researcher is able to construct $\succsim$. The agent considers X as the set of consequences and I again assume that the researcher knows X . What is missing is the state space. I make no assumption about whether or not the agent constructs a state space. However, I do assume that the researcher does not presuppose any state space and models the agent’s behavior as if there were no state space.

Without a state space, alternatives cannot be described as Savage acts. Hence, the obvious question is: What is the use of the set of consequences? The set of consequences, X , makes it possible to describe what the alternatives mean to the agent. Each alternative, f , is associated with a unique subset of X , X_f , and this subset is interpreted as the set of outcomes that the agent considers to be possible when she chooses f . When X_f is countable, there is no problem with understanding what “possible outcomes” means. In this case X_f is a collection of all consequences to which the agent assigns a non-zero probability. However, this term need not be clear when X_f is uncountable. Instead it is likely that X_f and, for some $x \in X_f$, $X_f \setminus \{x\}$ are two sets that are indistinguishable from the probabilistic perspective. Obviously this is a source of potential problems, since each alternative should be identified with a unique subset of X . (I address this issue in greater detail in Subsection 3.2.) I assume that the researcher knows how alternatives are associated with the subsets of X .

In formal terms, an alternative is a pair, (f, X_f) , where f is a label and X_f is a subset of X . I define alternatives in this way because it is possible that there are two alternatives that share the same set of consequences but which are not evaluated as being the same or equally good. Consider a bet on the next winner of the U.S. presidential elections. There are two alternatives: a Democrat, denoted by f_D , and a Republican, denoted by f_R . Suppose that picking the winner yields \$100, while picking the loser costs the agent \$50. Each alternative is associated with the same set of outcomes—namely $\{\$100, -\$50\}$. However, it need not be true that the agent perceives f_D and f_R as the same object or that the agent is indifferent between f_D and f_R .

As standard in the literature, the strict part of $\succsim$ is denoted by $\succ$ and is defined in the usual way: $f \succ g$ if and only if it is true that $f \succsim g$ and it is false that $g \succsim f$. I write $x \in X_f$ whenever x is a consequence that, at least in the agent’s mind, may obtain if she chooses f . Let f_x denote a singleton alternative that yields consequence x —that is, $X_{f_x} = \{x\}$. A singleton alternative generates no uncertainty—if the agent chooses

f_x , then she knows the final outcome will be x . Although not formally correct, I will abuse the notation and write x instead of f_x .

3.2 Set-up and objectives

Fix the set of alternatives, $\mathcal{F}$, the set of consequences, X , and the agent's preference relation, $\succsim$, defined over the set of alternatives. It is useful to think of the agent as facing a two-stage procedure: First, the agent chooses an alternative; second, Nature picks one element from the set of consequences attributed to that alternative. Since the agent has no control over Nature's choice, she assigns probability λ_f to the set of outcomes associated with f . The alternatives are evaluated by their subjective expected utility—in other words, if u is a utility function defined over the set of consequences, then f is preferred over g if and only if $\int_{X_f} u(x) d\lambda_f \geq \int_{X_g} u(x) d\lambda_g$.

It is important that such a probability measure, λ_f , assigned to alternative, f , captures the agent's perception of f . In the case of countable alternatives, it is clear: If x is considered to be possible when f is chosen, then λ_f assigns a non-zero weight to x . However, with uncountable X_f , it is not possible to require that $\lambda_f(x) > 0$ for each $x \in X_f$. This implies that there is no straightforward definition of the “possible outcomes” for uncountable alternatives. Suppose that $X = [0, 1]$ and take $X_f = [0, 0.5]$. If $\lambda_f(0) = 0$, then it is not clear whether X_f should be defined as $[0, 0.5]$ or $(0, 0.5]$. I solve this problem by implicitly assuming that the agent is indifferent between X_f and the closure of X_f and, in Assumption 3.1, I focus only on those uncountable alternatives that are closed in X . In consequence, λ_f is said to capture agent's perception about an alternative whenever the support of λ_f is X_f . (Note that if λ is a probability measure defined on a metric space, X , then the support of λ is defined as the smallest closed subset of X , C , such that $\lambda(C) = 1$.) However, without additional requirements, it is not necessarily true that for a given X_f there exists a probability measure with support X_f . This issue is also resolved in Assumption 3.1, which requires that uncountable X is a separable metric space. This guarantees that for a given closed subset of X , X_f , there exists a probability measures on X with support X_f (see Exercise 4, Chapter 7.1 in Dudley (2002)).

The idea behind a two-stage procedure discussed above—first, choice of the agent; second, choice of Nature—is derived from the theory of Savage. Recall that in a model with state space, each alternative is a function that is evaluated by its expected utility, $\int_{\Omega} u(f(\omega)) d\lambda$. Note that f and λ generate the probability measure on the set of consequences denoted as λ_f . Such a λ_f can be identified as the agent's subjective belief associated with f , and support of λ_f can be interpreted as the set of consequences that the agent considers to be possible if she chooses f .

To summarize, the model of agent's behavior consists of three elements—the set of alternatives, the preference relation defined over that set of alternatives, and the set of consequences—with alternatives described as sets of possible consequences. I impose the assumptions that allow for construction of (a) a utility function, u ,

defined over the set of consequences and (b) for each alternative, f , a probability measure, λ_f , with support X_f , such that f is evaluated by its subjective expected utility, $\int_{X_f} u(x)d\lambda_f$.

3.3 Utility representation

The first assumption, as previously indicated, is technical and guarantees that, for each $X_f \subset X$, there exists a probability measure, λ_f , such that the support of λ_f is X_f .

Assumption 3.1 *Set of consequences.*

If X is uncountable, then X is a separable metric space endowed with the Borel σ -algebra. In this case, if X_f is uncountable, then X_f is a closed subset of X and is endowed with the Borel σ -algebra.

Assumption 3.2 is an axiom of choice theory and guarantees the existence of utility function, V , representing the agent's preference relation.

Assumption 3.2 *Preference relation.*

$\succsim$ is a complete and transitive binary relation on $\mathcal{F}$. If $\mathcal{F}$ is uncountable, then $\succsim$ is continuous.

In the case of an uncountable set of alternatives, combining Assumption 3.2 with Assumption 3.3a guarantees that V is a continuous function (see Debreu (1954, 1964)). Assumption 3.3b requires that the set of alternatives be rich enough to allow for the derivation of a utility function, u , defined on the set of consequences, X . Such a function is defined in a natural way as $u(x) := V(f_x)$.

Assumption 3.3 *Set of alternatives.*

- (a) *If $\mathcal{F}$ is uncountable, then $\mathcal{F}$ is a separable metric space.*
- (b) *For each $x \in X$, there is $f_x \in \mathcal{F}$.*

Let $\max f$ and $\min f$ denote the maximum and minimum (with respect to $\succsim$) elements of X_f , respectively. That is, for each $x \in X_f$, $\max f \succsim x \succsim \min f$. Assumption 3.4 gives the key requirement.

Assumption 3.4 *Max–Min boundaries.*

- (a) *For each $f \in \mathcal{F}$, there exist $\max f$ and $\min f$.*
- (b) *For each $f \in \mathcal{F}$, if $\max f \succ \min f$, then $\max f \succ f \succ \min f$.*

Note: Instead of assuming the existence of $\max f$ and $\min f$, one can require that X be a compact metric space (Assumption 3.1). In that case, X_f , a closed subset of X , is also compact, and since u is a continuous function, the Weierstrass Theorem dictates that both $\max f$ and $\min f$ exist.

Assumption 3.4b can be seen as an extension of the Simple Dominance Axiom introduced by Barberà (1977), which holds that $\{x\} \succ \{y\}$ implies $\{x\} \succ \{x, y\} \succ \{y\}$.

It is crucial that Assumption 3.4 be formulated in terms of $\succ$ instead of $\succsim$. When f is not a singleton, then such an act should not be indifferent to its best or worst element. Suppose that $X_f = \{x_1, x_2, x_3\}$ where $x_1 \succ x_2 \succ x_3$. If $f \sim x_1$, then λ_f must be a degenerate measure that assigns 1 to x_1 . However, this would violate the fact that the agent considers all three consequences to be possible.

In order to better understand the limits imposed by Assumption 3.4, consider two possible violations of Assumption 3.4. First, suppose that the agent prefers $\min f$ to f . That agent does not want to try to improve her well-being, even though there are no costs involved in attempting this. After all, by choosing f , the worst possible result is $\min f$. Next consider the agent who prefers f to $\max f$. That is, instead of having a certainty that she will get the best possible outcome that f can generate, $\max f$, the agent picks f . Although Assumption 3.4 may appear to be an obvious and harmless requirement, I assert that, when appropriate, it is reasonable to expect that Assumption 3.4 will be violated. Consider the agent's problem of choosing the destination for her next summer vacation. For simplicity, Ann chooses from the set of U.S. states, $\mathcal{F}$. Let X be the set of consequences representing Ann's different levels of satisfaction derived from summer trip: "very good," "good," "acceptable," "bad," and "very bad." If h represents Hawaii, then Ann, who has already visited the Hawaiian islands, identifies h with "very good," her best possible consequence. This year, Ann decides to go to Georgia, g , which she has never visited before. Ann believes that Georgia will yield one of three consequences, "very good," "good," or "acceptable." Ann strictly prefers Georgia over Hawaii and her behavior violates Assumption 3.4. This choice may appear to be puzzling. However, visiting Georgia means not only having a vacation but also learning what a vacation in Georgia means. That is, choosing a location she has never visited before yields the additional benefit of learning—that is, after Ann experiences Georgia, g , she learns what it means to go there, and g become a singleton. Ann's utility over $\mathcal{F}$ is a weighted sum, $V(f) = EU(f) + \alpha|f|$, where $EU(f)$ is the expected utility derived from f and $|f|$ is a cardinality of f . (In this context, cardinality can be interpreted as a measure of how well the agent knows f . For Hawaii, $|h| = 1$, but for Georgia, $|g| = 3$.) A parameter, α , measures the importance of learning. Although a trip to Georgia yields expected value that cannot be bigger than the (certain) value generated by the Hawaii vacation, Ann chooses Georgia over Hawaii. This happens because α is strictly bigger than zero (in other words, Ann likes to learn). Alternatively, consider Bob who spends every vacation in Texas, t , and considers t to be "good." Bob knows that there are other alternatives—in particular, California, c , which he perceives to be "very good" or "good." Going to California will guarantee an expected value that is at least the same as the trip to Texas. However, Bob's α is negative—that is, he dislikes learning—and he prefers to stick to what he already knows best. In consequence, he, like Ann, violates Assumption 3.4.

Assumptions 3.1 through 3.4 imply that the agent behaves as if she were evaluating the alternatives by their subjective expected utility.

Lemma 3.1

Fix X , $\mathcal{F}$, and $\succsim$. If Assumptions 3.1, 3.2, 3.3, and 3.4 hold, then:

1. *for each $f \in \mathcal{F}$ there is a probability measure, λ_f , defined on X such that the support of λ_f is X_f ; and*
2. *there exist bounded and continuous functions, $u : X \rightarrow \mathbb{R}$ and $V : \mathcal{F} \rightarrow \mathbb{R}$, such that V represents $\succsim$ and is defined as $V(f) := \int_{X_f} u(x) d\lambda_f$.*

4 Step 2: State space and belief over state space

Lemma 3.1 shows that the agent behaves as if she were ranking alternatives according to the subjective expected utility they yield. Note that there still are no state space, Ω , and no probability, λ , defined on that state space, to translate the structure presented in Section 3 into an SEU model. These are provided next in Proposition 4.1.

Proposition 4.1

Fix X , $\mathcal{F}$, and $\succsim$. Suppose that Assumptions 3.1, 3.2, 3.3, and 3.4 hold, and assume the Axiom of Choice. Then:

1. *there exists a state space, Ω , such that each $f \in \mathcal{F}$ is representable as a measurable function $f : \Omega \rightarrow X$; and*
2. *there exist bounded and continuous functions, $u : X \rightarrow \mathbb{R}$ and $V : \mathcal{F} \rightarrow \mathbb{R}$, and a probability measure, λ , defined on Ω such that V represents $\succsim$ and is defined as $V(f) := \int_{\Omega} u(f(\omega)) d\lambda$.*

In Proposition 4.1, the state space is defined as the product of sets of consequences associated with the alternatives, $\Omega := \times_{f \in \mathcal{F}} X_f$. As such, each state specifies the outcome of each alternative. This is not a novel idea (see subsection 11.1 in Gilboa (2009)). One important objection to Proposition 4.1 concerns the construction of belief, λ , which is defined as a product of the probabilities associated with alternatives, λ_f s. As such, the alternatives are assumed to be outcome-independent—that is, knowing the realization of alternative f does not change the agent’s evaluation of alternative g . In general, this assumption need not be correct. However, in order to determine the agent’s perception of the relationships among alternatives and derive a measure over Ω that is not a product measure of λ_f s, it is necessary to introduce more complicated alternatives, leaving this as an open problem.

The construction of Ω and λ in Proposition 4.1 provides the basis for the SEU-rationalization of the Ellsberg experiment proposed in Table 1. Let X_r stand for the set of possible consequences when the agent bets on the red ball, f_r . Similarly, there are X_g , X_{rb} , X_{gb} . Note that $X_r = X_g = X_{rb} = X_{gb} = \{S, F\}$. Following Lemma 3.1, I derive λ_f s for the four acts in question: $\lambda_r(S) = \frac{4}{6}$, $\lambda_g(S) = \frac{2}{6}$, $\lambda_{rb}(S) = \frac{5}{6}$, and $\lambda_{gb}(S) = \frac{3}{6}$. As suggested by Proposition 4.1, the state space is built as $\Omega = X_r \times X_g \times X_{rb} \times X_{gb}$, and agent's belief, λ , is defined as a product measure. As noted in Section 2, Table 1 provides an SEU-rationalization of the Ellsberg experiment, but at the cost of assigning high probability to irrational states. That construction from Table 1 then leads to the creation of an SEU model with ε -small irrationality in Table 2.

5 Related literature

This paper belongs to a large body of literature that is concerned with an agent ranking sets of objects. (See Barberà et al. (2004) for the most complete review of this literature.) The models in this literature consist of two stages. In the first stage, the agent chooses a subset of X . What happens at the second stage divides these models into the following two groups:

- A. models in which, at the second stage, Nature picks the final outcome. (See Section 3 in Barberà et al. (2004).) These models are known as the models of choice under “complete uncertainty,” have the same structure as the model in Section 3, and are comparable to Lemma 3.1. However, because they are not concerned with construction of the agent's state space, they are not comparable to Proposition 4.1.
- B. models in which, at the second stage, the agent picks the final outcome. (See Section 4 in Barberà et al. (2004).) Here the non-existence of Nature implies that these models are not comparable to Lemma 3.1. However, the models of “preference for flexibility,” which belong to this group, also aim to construct the state space from the agent's preference relation and, as such, are comparable to Proposition 4.1.

The models in both groups share one common feature that separates them from the set-up I propose. In these models, the agent's preference relation is defined over the power set of X . Such a construction implicitly assumes that two alternatives with the same set of consequences should be perceived as the same object or, at least, that the agent should be indifferent between them. As the example of betting on the winner of the next U.S. presidential election indicates (Subsection 3.1), this assumption is too restrictive for the objective of this paper. One possible solution is to define the consequences in a different way. For instance, instead of $x_1 = \$100$, one can think of two consequences: $y_1 = \text{“win \$100 when a Democrat wins”}$ and $y_2 = \text{“win \$100 when a Republican wins”}$. Similarly, $x_2 = \$ - 50$ can be replaced by $y_3 = \text{“lose \$50 when a Democrat loses”}$ and $y_4 = \text{“lose \$50 when a Republican loses.”}$ With this new set of consequences, betting that the winner will be a Democrat or a Republican can be identified by $X_D = \{y_1, y_3\}$ and $X_R = \{y_2, y_4\}$, respectively. Clearly these two bets differ in terms of the sets of outcomes. However, another problem

emerges. With the set of consequences $X = \{y_1, y_2, y_3, y_4\}$, it is possible that the agent will be unable to express her opinion about all of the subsets of X . In particular, singletons of X do not provide a complete picture of what the agent might win or lose. For instance, $X_A = \{y_1\}$ indicates that the agent wins \$100 if a Democrat wins but does not address a Republican victory. If the agent is unable to rank X_A , then it is not possible to construct a utility function over X and, in consequence, carry out SEU-rationalization. There are also subsets of X that seem contradictory. For example, consider $X_B = \{y_1, y_4\}$, which indicates that the agent both gains \$100 when a Democrat wins and loses \$50 when a Democrat wins (i.e., when a Republican loses). In order to avoid these problems, I define an alternative as a pair consisting of a name, f , and a set of consequences, X_f .

The essential difference between the models in Group A and Lemma 3.1 is the use of probabilities. As stated by Barberà et al. (2004), “[T]he term ‘complete uncertainty’ refers to a situation where the agent knows the set of possible consequences of an action but cannot assign probabilities to those outcomes.” My main objective is to model the agent’s choice *as if* it were derived from constructing subjective beliefs for alternatives and evaluating alternatives according to their subjective expected utilities. However, this paper is not the first to assume that the agent behaves as if she were constructing beliefs over the sets of possible outcomes. Ahn (2008), for instance, starts with the finite set X and the set of probability measures on X , $P(X)$. In his model, agent’s preference relation is defined for singletons and regular subsets of $P(X)$. (A is a regular subset if the closure of the interior of A is the same as A .) Let v denote the continuous utility function defined on $P(X)$, which does not necessarily take the expected utility form. Ahn (2008) derives a probability measure, μ , on $P(X)$ and a representation of the form $V(A) = v(I)$ if $A = \{I\}$, and $V(A) = \frac{\int_A v d\mu}{\mu(A)}$ if A is not a singleton. In addition, in Gravel et al. (2007), there is a set of consequences, X , and the agent’s preference relation is defined over the set of all finite subsets of X . The authors obtain a representation, $V(A) = \frac{\sum_{a \in A} u(a)}{|A|}$, where u is a utility function on X and $|A|$ denotes the cardinality of A . Dividing $\sum_{a \in A} u(a)$ by the cardinality of A can yield the interpretation that the agent believes that each element of A has the same probability of occurring. Hence, this representation is called the Uniform Expected Utility.

Within Group A, there is a stream of models that, as with Hurwicz’s α -criterion,¹ develop the idea of evaluating sets by taking only their best and worst outcomes into consideration.² The assumptions used in Lemma 3.1 imply an alternative representation that is in accordance with that idea. Given V , there is a function $\alpha : \mathcal{F} \rightarrow (0, 1)$ such that

$$V(f) = \alpha(f) \cdot u(\max f) + (1 - \alpha(f)) \cdot u(\min f). \quad (4)$$

¹See Arrow and Hurwicz (1972), Hurwicz (1951), and Milnor (1954).

²See Arlegi (2003), Arlegi et al. (2005), Barberà et al. (1984), Barberà and Pattanaik (1984), Larbi et al. (2010), Bossert (1989), Bossert et al. (1994), Bossert et al. (2000), Dutta and Sen (1996), Nehring and Puppe (1996), Olszewski (2007), Pattanaik and Xu (1998), and Pattanaik and Xu (2000).

By Assumption 3.4, $u(\min f) < V(f) < u(\max f)$, and hence $\alpha(f)$ is defined as $\frac{u(\max f) - V(f)}{u(\max f) - u(\min f)}$. In (4), α can be interpreted as a probability measure on X_f , which assigns weights $\alpha(f)$ and $1 - \alpha(f)$ to the best and the worst consequence of f , respectively. This functional representation resembles the result obtained by Ghirardato (2001). In his model, the state space, Ω , is fixed, and the agent perceives acts as set-valued functions. That is, $f(\omega)$ is not required to be a singleton. In fact, (4) can be considered to be an extreme version of Ghirardato (2001), in which the agent perceives no state space. The representation in (4) is not appropriate for this paper, since it does not capture the agent's perception of alternatives. Recall that X_f represents all outcomes that are deemed possible by the agent when she chooses f . For that reason, I focus this paper on the probability measures with support X_f . In (4), however, the support of α is not X_f but, rather, $\{\max f, \min f\}$.

In Group B, we look to the literature on “preference for flexibility” that begins with Kreps (1979)³ and is concerned with the agent who is not sure about her future preferences. This uncertainty is captured by a state space, S , in which a state, s , represents a possible future preference relation over set X . My paper implicitly assumes that the agent does not face uncertainty regarding her future self and consequently the construction of state space differs from the construction proposed in Kreps (1979). In Kreps (1979), the agent evaluates a subset of X , A , by $V(A) = \sum_{s \in S} \max_{x \in A} u(s, x)$, where $u(s, x)$ is an ex-post, state-dependent utility function. In my paper, however, a menu is evaluated by its subjective expected utility, and utility over X is state independent. Kreps (1992) reinterprets the original “preference for flexibility” model as representing the agent who faces unforeseen contingencies—that is, the agent who is unable to account for all possible future uncertainties. In my paper, however, the model does not represent the agent who does not have a complete picture of uncertainty but, rather, the researcher who does not know which uncertainties the agent takes into account.

6 Final remarks

In this paper, I consider the problem of the researcher who seeks to SEU-rationalize the agent's behavior. The researcher starts with the set of alternatives, the set of consequences, and the preference relation. He builds a state space and belief over that state space such that the agent's choices appear to be in accordance with the SEU theory. Of course, not knowing the agent's state space is only one of several problems that the researcher may face.

To make the discussion more precise, let $(\mathcal{F}, \Omega, X, P)$ be called a model. As previously noted, $\mathcal{F}$, Ω , and X denote the set of alternatives, the state space, and the set of consequences, respectively. Recall that elements of the set of alternatives are just labels. P translates these labels into meaningful objects. If the state space

³See also Barberà and Grodal (2011), Chatterjee and Krishna (2011), Dekel et al. (2001), Dekel et al. (2007), Epstein et al. (2007), Epstein and Seo (2009), Nehring (1999), Ozdenoren (2002), Sadowski (2010), Sagi (2006), and Sarver (2008).

is assumed, then P translates alternatives into acts (i.e., the functions from Ω to X)—that is, $P : \mathcal{F} \rightarrow X^\Omega$. If, as in this paper, the state space is the empty set, then P translates the alternatives into the subsets of X —that is, $P : \mathcal{F} \rightarrow 2^X$.

Next fix an agent who analyzes some problem and let $(\mathcal{F}^A, \Omega^A, X^A, P^A)$ be her model. There also is a researcher who analyzes the agent, and his model is $(\mathcal{F}^R, \Omega^R, X^R, P^R)$. In general, the models of the agent and the researcher can differ. If this is the case, the researcher faces the risk of drawing incorrect conclusions. Since a model consists of four elements, there are four possible differences between the agent’s and researcher’s models which I discuss below.

1. First, it is possible that the set of alternatives that the agent considers is not the same as the researcher’s understanding of what the set of alternatives is—that is, $\mathcal{F}^A \neq \mathcal{F}^R$. This problem is not specific to the choice theory under uncertainty but rather belongs to the literature on choice theory under certainty. In the case of $\mathcal{F}^A \subset \mathcal{F}^R$, this issue is analyzed by Masatlioglu et al. (2011), but analysis of the opposite case, $\mathcal{F}^R \subset \mathcal{F}^A$, remains an open problem.
2. Second, the agent’s state space might differ from the researcher’s state space—that is, $\Omega^A \neq \Omega^R$. In general, the researcher’s data set is not equipped with the agent’s state space. If the researcher chooses to use a state space in his analysis, he must assume some specific state space, thereby, and risks making an incorrect assumption. Constructing a state space is not a trivial task, since there generally is no obvious way to approach this issue. Moreover, even if the problem appears to be simple, there could be more than one way to construct a reasonable state space. In the business literature, there is a concept of “scenario planning” (see Schoemaker (1995)) that, essentially, attempts to create a tool that helps managers construct states. But even this method does not guarantee that two different managers would come up with the same, or similar, state spaces.
3. Third, the agent and researcher might consider distinct sets of consequences—that is, $X^A \neq X^R$. This problem is similar to the agent and researcher having different state spaces, in the sense that both the state space and the set of consequences are unobservable elements of the agent’s decision-making process. The analysis of this issue remains an open problem.
4. Fourth, the agent’s translation of alternatives into acts, or subsets of X , might differ from how the researcher perceives the alternatives—that is, $P^A \neq P^R$. Suppose, for simplicity, that the agent and researcher share the same state space, Ω , and the same set of consequences, X . It is possible that a given alternative, f , need not mean the same thing to both the agent and the researcher. In other words, there could be a state, ω , in which the consequence of choosing f at that state, $f(\omega)$, is x_1 in the eyes of the agent and x_2 in the eyes of the researcher, with $x_1 \neq x_2$. However, this issue belongs to the literature on choice theory under certainty and remains an open problem.

In this paper, I presume that $\mathcal{F}^A = \mathcal{F}^R$, $X^A = X^R$, and $P^A = P^R$. I begin with $\Omega^R = \emptyset$, in order to

derive a state space, Ω , and a belief, λ , that SEU-rationalize the agent's behavior. It also is possible that the researcher begins with some specific state space, $\Omega^R \neq \emptyset$, and ponders whether it is possible to enlarge or reduce such a state space in order to SEU-rationalize the agent's behavior. This also remains an open problem.

Appendix

Proof of Lemma 3.1

As already explained, Assumptions 3.2 and 3.3 guarantee the existence of continuous and bounded V and u . However, it remains to be proven that (a) for each alternative there exists a probability measure, λ_f , with support X_f and (b) that $V(f) = \int_X u(x)d\lambda_f$.

Let $P(X)$ denote the collection of probability measures on X that are Borel for uncountable X . Fix f and let $P_f(X)$ be a subset of $P(X)$ that consists of probability measures on X with support X_f . For countable X_f , the non-emptiness of $P_f(X)$ is obvious. For uncountable X_f , the non-emptiness is guaranteed by Assumption 3.1. Let $\Psi : P(X) \rightarrow \mathbb{R}$ be the integral function $\Psi(\lambda) := \int_X u d\lambda$. The remaining part of proof shows that for a given f there exists $\lambda_f \in P_f(X)$ such that $\Psi(\lambda_f) = V(f)$.

Consider an alternative, f , such that $\max f \sim \min f$. However, Assumption 3.4 does not apply to such an f , since the agent is indifferent among the elements of X_f and consequently u is constant over X_f , so any λ_f with support X_f would do the job. Hence, consider an act, f , such that $\max f \succ \min f$. The idea of the proof is to find two probability measures, λ_1 and λ_2 , that both belonging to $P_f(X)$ such that $c_1 = \Psi(\lambda_1) < V(f) < \Psi(\lambda_2) = c_2$. With c_1 and c_2 in hand, it is possible to find the unique α such that $\alpha c_1 + (1 - \alpha)c_2 = V(f)$. With this α , the desired λ_f is defined as $\lambda_f = \alpha\lambda_1 + (1 - \alpha)\lambda_2$.

It remains to be proven that, for a given f , such probability measures, λ_1 and λ_2 , indeed exist. Note that due to Assumption 3.4, $\Psi(\delta_{\min f}) < V(f) < \Psi(\delta_{\max f})$, where δ_x is a degenerate probability measure with mass 1 at x . One way to show that there always is λ_1 such that $\Psi(\lambda_1) < V(f)$ involves proving that there is a sequence of probability measures, $\{\lambda_1^n\}$, such that each $\lambda_1^n \in P_f(X)$ and $\Psi(\lambda_1^n) \rightarrow \Psi(\delta_{\min f})$. Existence of such a sequence guarantees that there is some N such that $\Psi(\delta_{\min f}) < \Psi(\lambda_1^N) < V(f)$. Such a λ_1^N then becomes the desired λ_1 . Since the existence of λ_2 is proven in the same way, I focus on λ_1 . Consider a countable X_f . I construct $\{\lambda_1^n\}$ in the following way. Let $\lambda_1^n(\min f) = 1 - \frac{1}{n}$ and let $\sum_{x \in X_f \setminus \{\min f\}} \lambda_1^n(x) = \frac{1}{n}$. If X_f is finite with cardinality K , then let $\lambda_1^n(x) = \frac{1}{(K-1)n}$ for $x \in X_f \setminus \{\min f\}$. If X_f is infinite, then let $X_f = \{x, x_0, x_1, \dots\}$, where $x = \min f$. For x_m , define $\lambda_1^n(x_m) = \frac{1}{2^n} \cdot \frac{1}{2^m}$. Note that $\sum_{m=0}^{\infty} \lambda_1^n(x_m)$ is a convergent geometric series such that $\sum_{m=0}^{\infty} \frac{1}{2^n} \cdot \frac{1}{2^m} = \frac{1}{2^n}$. With $\{\lambda_1^n\}$ in hand, it is then possible to compute $\Psi(\lambda_1^n) = (\sum_{x \in X_f} u(x)\lambda_1^n(x))$ and observe that the sum converges, as desired, to $\Psi(\delta_{\min f})$ as n goes to infinity. Finally, consider an uncountable X_f . As already noted, the fact that X is separable and X_f is a

closed subset of X implies the existence of probability measure $\hat{\lambda}_f$ on X with the support X_f . Let λ_1^n be a weighted measure that assigns $\frac{1}{n}$ to $\hat{\lambda}_f$ and $1 - \frac{1}{n}$ to $\delta_{\min f}$. Each λ_1^n has support X_f and $\Psi(\lambda_1^n) \rightarrow \Psi(\delta_{\min f})$ when $n \rightarrow \infty$. ■

Proof of Proposition 4.1

By the Axiom of Choice, for an arbitrary $\mathcal{F}$, there exists $\Omega := \times_{f \in \mathcal{F}} X_f$. An alternative, f , is transformed into an act by defining $f : \Omega \rightarrow X$ as a projection function. As such, function f is measurable. Let τ be a finite nonempty subset of $\mathcal{F}$. Let $\Omega_\tau := \times_{f \in \tau} X_f$. Let λ_τ be the probability measure on Ω_τ constructed as a product of measures $\{\lambda_f\}$. By Proposition V.1.2 in Neveu (1965), there exists the unique probability measure, λ , on Ω that agrees with each λ_τ on cylinders. Hence, there is a state space, Ω , with associated probability, λ . The final step requires establishing that $V(f)$ from Lemma 3.1 is equal to $\int_\Omega u(f(\omega)) d\lambda$. Let Ω_{-f} denote the product $\times_{f \in \mathcal{F} \setminus \{f\}} X_f$. Note that $\Omega = \Omega_{-f} \times X_f$. By construction of λ , $\int_{\Omega_{-f} \times X_f} u(f(\omega)) d\lambda = \int_{X_f} u(x) d\lambda_f$. ■

References

- AHN, D. (2008): “Ambiguity Without a State Space,” *Review of Economic Studies*, 75, 3–28.
- ANSCOMBE, F. J. AND R. AUMANN (1963): “A Definition of Subjective Probability,” *Annals of Mathematical Statistics*, 34, 199–205.
- ARLEGI, R. (2003): “A note on Bossert, Pattanaik and Xu’s Choice under complete uncertainty: axiomatic characterization of some decision rules,” *Economic Theory*, 22, 219–225.
- ARLEGI, R., M. BESADA, J. NIETO, AND C. VÁZQUEZ (2005): “A note on Bossert, Pattanaik and Xu’s Choice under complete uncertainty: axiomatic characterization of some decision rules,” *Mathematical Social Sciences*, 49, 1–15.
- ARROW, K. J. AND L. HURWICZ (1972): “An Optimality Criterion for Decision-Making under Ignorance,” in *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*, ed. by C. Carter and J. Ford, Oxford: Basil Blackwell.
- BARBERÀ, S. (1977): “The Manipulation of Social Choice Mechanisms that Do Not Leave “Too Much” to Chance,” *Econometrica*, 45, 1573–1588.
- BARBERÀ, S., C. R. BARRETT, AND P. K. PATTANAİK (1984): “On Some Axioms for Ranking Sets of Alternatives,” *Journal of Economic Theory*, 33, 301–308.

- (2004): “Ranking Sets of Objects,” in *Handbook of Utility Theory Vol. 2*, ed. by P. Hammond and C. Seidl, Springer.
- BARBERÀ, S. AND B. GRODAL (2011): “Preference for flexibility and the opportunities of choice,” *Journal of Mathematical Economics*, 47, 272–278.
- BARBERÀ, S. AND P. K. PATTANAIK (1984): “Extending an Order on a Set to the Power Set: Some Remarks on Kannai and Peleg’s Approach,” *Journal of Economic Theory*, 32, 185–191.
- BOSSERT, W. (1989): “On the Extension of Preferences over a Set to the Power: An Axiomatic Characterization of a Quasi-ordering,” *Journal of Economic Theory*, 49, 84–92.
- BOSSERT, W., P. K. PATTANAIK, AND Y. XU (1994): “Ranking Opportunity Sets: An Axiomatic Approach,” *Journal of Economic Theory*, 63, 326–345.
- (2000): “Choice under Complete Uncertainty: Axiomatic Characterizations of Some Decision Rules,” *Economic Theory*, 16, 295–312.
- BREWER, K. R. W. (1963): “Decisions under Uncertainty: Comment,” *Quarterly Journal of Economics*, 77, 159–161.
- CHATTERJEE, K. AND V. R. KRISHNA (2011): “On Preferences with Infinitely Many Subjective States,” *Economic Theory*, 46, 85–98.
- DEBREU, G. (1954): “Representation of a Preference Ordering by a Numerical Function,” in *Decision Processes*, ed. by R. Thrall, C. Coombs, and R. Davis, Wiley.
- (1964): “Continuity Properties of Paretian Utility,” *International Economic Review*, 5, 285–293.
- DEKEL, E., B. L. LIPMAN, AND A. RUSTICHINI (2001): “Representing Preferences with a Unique Subjective State Space,” *Econometrica*, 69, 891–934.
- DEKEL, E., B. L. LIPMAN, A. RUSTICHINI, AND T. SARVER (2007): “Representing Preferences with a Unique Subjective State Space: A Corrigendum,” *Econometrica*, 75, 591–600.
- DUDLEY, R. M. (2002): *Real Analysis and Probability*, Cambridge University Press.
- DUTTA, B. AND A. SEN (1996): “Ranking Opportunity Sets and Arrow Impossibility Theorems: Correspondence Results,” *Journal of Economic Theory*, 71, 90–101.
- EICHBERGER, J. AND D. KELSEY (2009): “Ambiguity,” in *The Oxford Handbook of Rational and Social Choice*, ed. by P. Anand, P. Pattanaik, and C. Puppe, Oxford University Press.
- ELLSBERG, D. (1961): “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 75, 643–669.

- EPSTEIN, L. G. (2010): “A Paradox for the “Smooth Ambiguity” Model of Preference,” *Theoretical Economics*, 78, 2085–2099.
- EPSTEIN, L. G., M. MARINACCI, AND K. SEO (2007): “Coarse Contingencies and Ambiguity,” *Theoretical Economics*, 2, 355–394.
- EPSTEIN, L. G. AND K. SEO (2009): “Subjective States: A More Robust Model,” *Games and Economic Behavior*, 67, 408–427.
- FRISCH, D. AND J. BARON (1988): “Ambiguity and rationality,” *Journal of Behavioral Decision Making*, 1, 149–157.
- GHIRARDATO, P. (2001): “Coping with Ignorance: Unforeseen Contingencies and Non-additive Uncertainty,” *Economic Theory*, 17, 247–276.
- (2010): “Ambiguity,” in *Encyclopaedia of Quantitative Finance*, ed. by R. Cont, J. Wiley and Sons.
- GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): “Differentiating Ambiguity and Ambiguity Attitude,” *Journal of Economic Theory*, 118, 133–173.
- GILBOA, I. (2009): *Theory of Decision under Uncertainty*, Cambridge University Press.
- GILBOA, I., F. MACCHERONI, M. MARINACCI, AND D. SCHMEIDLER (2010): “Objective and Subjective Rationality in a Multiple Prior Model,” *Econometrica*, 78, 755–770.
- GILBOA, I. AND L. SAMUELSON (2011): “Subjectivity in Inductive Inference,” *Theoretical Economics*, forthcoming.
- GILBOA, I. AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- (1994): “Additive representations of non-additive measures and the Choquet integral,” *Annals of Operations Research*, 52, 43–65.
- GRAVEL, N., T. MARCHANT, AND A. SEN (2007): “Ranking Completely Uncertain Decisions by the Uniform Expected Utility Criterion,” *IDEP Working Papers*, 750.
- HANSEN, L. P. AND T. J. SARGENT (2001): “Robust Control and Model Uncertainty,” *American Economic Review*, 91, 60–66.
- HEY, J. D., G. LOTITO, AND A. MAFFIOLETTI (2010): “The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity,” *Journal of Risk and Uncertainty*, 41, 81–111.
- HURWICZ, L. (1951): “Optimality Criteria for Decision Making Under Ignorance,” *Cowles Commission Discussion Paper: Statistics*, 370.

- KADANE, J. B. (1992): “Healthy scepticism as an expected-utility explanation of the phenomena of Allais and Ellsberg,” *Theory and Decision*, 32, 57–64.
- KARNI, E. (2008): “Unknowable states and choice-based definitions of subjective probabilities,” *Economics Letters*, 99, 534–536.
- KREPS, D. M. (1979): “A Representation Theorem for “Preference for Flexibility”,” *Econometrica*, 47, 565–577.
- (1992): “Static Choice and Unforeseen Contingencies,” in *Economics Analysis of Markets and Games: Essays in Honor of Frank Hahn*, ed. by P. Dasgupta, D. Gale, O. Hart, and E. Maskin, The MIT Press.
- KÜHBERGER, A. AND J. PERNER (2003): “The Role of Competition and Knowledge in the Ellsberg Task,” *Journal of Behavioral Decision Making*, 16, 181–191.
- LARBI, R. B., RAMZI, S. KONIECZNY, AND P. MARQUIS (2010): “A Characterization of Optimality Criteria for Decision Making under Complete Ignorance,” *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning, Toronto*, 172–181.
- MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, 74, 1447–1498.
- MACHINA, M. J. (2003): “States of the World and the State of Decision Theory,” in *The Economics of Risk*, ed. by D. Meyer, W.E. Upjohn Institute for Employment Research.
- MASATLIOGLU, Y., D. NAKAJIMA, AND E. OZBAY (2011): “Revealed Attention,” *American Economic Review*, forthcoming.
- MILNOR, J. W. (1954): “Games against nature,” in *Decision Processes*, ed. by R. Thrall, C. Coombs, and R. Davis, Wiley.
- MORRIS, S. (1997): “Risk, Uncertainty, and Hidden Information,” *Theory and Decision*, 42, 235–270.
- MUKERJI, S. (1997): “Understanding the nonadditive probability decision model,” *Economic Theory*, 9, 23–46.
- MUKERJI, S. AND H. S. SHIN (2002): “Equilibrium Departures from Common Knowledge in Games with Non-Additive Expected Utility,” *Advances in Theoretical Economics*, 2, 1534–5963.
- NEHRING, K. (1999): “Preference for Flexibility in a Savage Framework,” *Econometrica*, 67, 101–119.
- NEHRING, K. AND C. PUPPE (1996): “Continuous Extension of an Order on a Set to the Power Set,” *Journal of Economic Theory*, 68, 456–479.
- NEVEU, J. (1965): *Mathematical Foundations of the Calculus of Probability*, Holden-Day.

- OLSZEWSKI, W. (2007): "Preferences Over Sets of Lotteries," *Review of Economic Studies*, 74, 567–595.
- OZDENOREN, E. (2002): "Completing the State Space with Subjective States," *Journal of Economic Theory*, 105, 531–539.
- PATTANAIK, P. K. AND Y. XU (1998): "On Preference and Freedom," *Theory and Decision*, 44, 173–198.
- (2000): "On Ranking Opportunity Sets in Economic Environment," *Journal of Economic Theory*, 93, 48–71.
- PULFORD, B. D. (2009): "Is Luck on my Side? Optimism, Pessimism, and Ambiguity Aversion," *Quarterly Journal of Experimental Psychology*, 62, 1079–1087.
- SADOWSKI, P. (2010): "Contingent Preference for Flexibility: Eliciting Beliefs from Behavior," *Economic Research Initiatives at Duke Working Paper Series*, 83.
- SAGI, J. S. (2006): "What is an "Endogenous State Space"?" *Economic Theory*, 27, 305–320.
- SARVER, T. (2008): "Anticipating Regret: Why Fewer Options May be Better," *Econometrica*, 76, 263–305.
- SAVAGE, L. J. (1972): *The Foundations of Statistics*, Dover Publications.
- SCHMEIDLER, D. (1989): "Subjective Probability and Expected Utility Without Additivity," *Econometrica*, 57, 571–587.
- SCHOEMAKER, P. J. H. (1995): "Scenario Planning: A Tool for Strategic Thinking," *Sloan Management Review*, 36, 25–40.
- SEGAL, U. (1987): "The Ellsberg Paradox and Risk: An Anticipated Utility Approach," *International Economic Review*, 28, 175–202.
- SINISCALCHI, M. (2008): "Ambiguity and Ambiguity Aversion," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Palgrave Macmillan, second ed.
- (2009): "Vector Expected Utility and Attitudes Towards Variation," *Econometrica*, 77, 801–855.
- STRZALECKI, T. (2010): "Axiomatic Foundations of Multiplier Preferences," *Econometrica*, 79.
- TRAUTMANN, S. T., F. M. VIEIDER, AND P. P. WAKKER (2008): "Causes of ambiguity aversion: Known versus unknown preferences," *Journal of Risk and Uncertainty*, 36, 225–243.
- WAKKER, P. P. (2011): "Bibliography," <http://people.few.eur.nl/wakker/refs/webrfrncs.doc>.